

**PROVISIONING OF SERVICES BASED ON DECLARATIVE DESCRIPTIONS OF A
RESOURCE STRUCTURE OF A SERVICE**

CROSS REFERENCES

The present application is cross-referenced to application entitled "Automatic Provisioning of Services Based on a High Level Description and an Infrastructure Description," having docket number YOR920040003US1, even dated herewith, and which is included herein by reference in entirety for all purposes.

FIELD OF THE INVENTION

The present invention is directed to provisioning and management of a set of computing services in a computing utility system, based on a declarative description of the resource structure of the desired service.

BACKGROUND

The cost and complexity of managing IT infrastructure continues to grow rapidly. Several factors contribute to this trend. First, IT infrastructures today are based on a distributed network of heterogeneous platforms and applications. In such an environment, resources, their capabilities, and behavior, are represented differently. They are therefore harder to compare and reason about. Interdependencies between resources, in which one resource's behavior is affected by another one, are not well represented or understood. Administrative personnel need to exercise knowledge of every platform, application and network appliance used, as well of how

1 they can be assembled together for a particular purpose.

2 Second, in a highly competitive business environment, businesses must respond quickly to
3 market changes. Such changes may impose new requirements on the IT infrastructure, such as
4 supporting new computing services or applications, upgrading resources, incorporating new ones,
5 or changing the network structure. Realizing these changes is often a manual, tedious, and error
6 prone process. In particular, as configuration changes are made, new IT management processes
7 need to be defined, and existing processes may need to be updated.

8 Finally, service providers are moving towards an SLA-based service delivery model in which the
9 set of resources allocated to a customer is dynamically adjusted based on workload and
10 performance. Reconfiguring infrastructure resources dynamically in response to customer needs
11 demands prompt attention from administrative personnel increasing operational cost. Therefore,
12 a clear requirement of businesses today is to reduce the cost of maintaining an IT infrastructure
13 by reducing the overall complexity and the level of required human operation.

14 A common approach to addressing these challenges is incorporating automation into the
15 operation of the system. Common tasks such as adding a server to a computing service when the
16 load increases, or installing software on a server, are automated, thereby reducing human
17 involvement, the time to complete the task, and the probability for human errors. Workflows are
18 often used as a vehicle for automation because they are particularly well suited for coordinating
19 the execution of a set of activities that are long lived, tracking progress of activities, and
20 incorporating human interaction where necessary. *Provisioning engines*, including a workflow
21 engine, and some useful set of workflows organized in some structured way, are emerging as a
22 means to achieve the goal of reducing the cost through automation.

23 Automating the operation of the infrastructure, even by utilizing a provisioning engine, does not
24 fully address the aforementioned challenges. Automation procedures are often specific to a
25 particular infrastructure, computing environment, and service. When coding an automation
26 procedure it is impossible to predict all future changes in the service or infrastructure. Therefore

1 statically defined automation procedures are likely to require change. For example, changing
2 from a one-tier to a two-tier architecture, or adding resources with new capabilities, can require a
3 complex re-implementation of the automation procedures. This task is further complicated by
4 the many interdependencies between resources that are often implicit, by the combinatorial large
5 number of possible allocations and configurations of a given set of resources, by the variety of
6 possible computing services with different requirements, and by the many and rapidly evolving
7 types of hardware and software resources. Therefore there is a need to be able to describe a
8 computing service independent of a particular infrastructure, to describe the resources in a
9 service provider's infrastructure and their interdependencies, and to automatically generate the
10 instructions to provision and manage the service on the resources in the infrastructure.

11 **SUMMARY OF THE INVENTION**

12 Thus, this invention provides a process for performing provisioning given a declarative
13 description of a desired resource structure. A method is provided which determines how to
14 assemble the desired resource structure from the building blocks available in the infrastructure.
15 The method generates provisioning actions that affect the infrastructure to create a resource
16 structure that matches the declarative description given in a Concrete Resource Model.
17 Automatic generation of provisioning actions facilitates consistent implementations, and reduces
18 error. The instructions could be embodied in a form such as a workflow that would serve as input
19 to a provisioning engine. The method can be used to perform provisioning, including creating a
20 new service environment, destroying an existing one, modifying the combination of resources
21 allocated to a computing environment, modifying their configuration, or any combination of the
22 above.

23 In an example embodiment, the present invention includes a method for provisioning in a
24 computing utility infrastructure. The method comprising: obtaining a Concrete Resource Model
25 describing a desired resource structure; and using the Concrete Resource Model to generate at
26 least one provisioning action to create a matching resource structure in the computing utility

1 infrastructure.

2 **BRIEF DESCRIPTION OF THE DRAWINGS**

3 The foregoing and other objects, aspects and advantages will be better understood from the
4 following detailed description of a preferred embodiment of the invention with reference to the
5 drawings, in which:

6 **Fig. 1** The system in which the invention is used;

7 **Fig. 2** High level view of the management elements that are described in this invention;

8 **Fig. 3** shows elements of the knowledge subsystem;

9 **Fig. 4** shows an example of a Concrete Resource Model;

10 **Fig. 5** shows an example of a system structure with a scientific computing cluster service
11 environment;

12 **Fig. 6** shows an example of an execution of the Concrete Model Processing Engine; and

13 **Fig. 7** shows a Concrete Model Processing Engine process.

14 **DESCRIPTION OF THE INVENTION**

15 The invention provides methods, systems and apparatus for generating provisioning actions given
16 a declarative description of the desired computing service resource structure and a separate
17 description of the infrastructure elements. The declarative description is given in a Concrete
18 Resource Model. A method of the invention determines how to assemble a desired service
19 environment from the building blocks available in the infrastructure, or how to change the
20 composition of resources allocated to an existing service environment to meet new requirements.
21 The provisioning actions that are generated can be embodied in a form such as a workflow that
22 would serve as input to a provisioning engine.

1 The operating environment may possess any of the characteristics listed below. Although the
2 present invention is not dependent on these characteristics, the method is general enough to
3 handle such conditions and characteristics.

4 First, resources may be allocated to customers in combinations which are heterogeneous,
5 may be interdependent, and vary over time.

6 Second, the service environments provided to each customer may be different. For
7 example, one customer may be provided resources for a web site, and another for a
8 scientific computing cluster. Resource types, quantities, dependencies, and allocation
9 patterns will thus vary between customers.

10 Third, there can be multiple ways to construct a service environment from the resources
11 in a service provider's infrastructure. A customer may have preferences or requirements
12 for particular variations of a given service environment. A service provider may also
13 have operational constraints that dictate which variations are acceptable.

14 Fourth, the infrastructure varies between service providers. Further, for a given service
15 provider, the infrastructure varies over time. These variations can be a result of upgrades
16 or additions to the physical infrastructure.

17 The system in which the present invention is used is a distributed computer system which has
18 multiple computing resources interconnected via a network. A diagram of such a system, which
19 might be found in a hosting or data center, is shown in Figure 1. The computing resources in the
20 physical infrastructure include, but are not limited to, processors 101, storage 103, firewalls 105,
21 and software 107. The software can be operating systems, middleware or applications. In Figure
22 1 the available software has been preconfigured into Business Processes, Business Applications,
23 or Service Environments 107. The hardware resources are connected by a network 109 as
24 indicated by the grid of lines interconnecting all of these resources. This network may be
25 configure into one or more tiers, where each tier is separated by a router or firewall. Software

resources are assigned to physical resources by a management infrastructure. In this environment a subset of the resources are assigned to the management infrastructure. In Figure 1 these resources 111, 113, and 115 are indicated by a dotted line surrounding them. The resources assigned to the management infrastructure run the management software described in this invention. This software manages the rest of the resources. The resources used for the management infrastructure are not assigned to customers of the hosted environment. The remaining resources are assigned to customers as necessary. It is expected that customers 117 will primarily receive service by connecting to their resources through the Internet 119. However, customers can receive service if they are connected to their resources by any means, such as through a direct connection to a managed resource.

Figure 2 is a high level view of the elements that are described in this invention. They can be divided into three groups. The physical infrastructure 201 is the actual set of inter-related resources as described in Figure 1. The knowledge subsystem 237 includes a set of management entities that represent and control the resources. In particular, every resource is represented and controlled by a service with some canonical interfaces, termed a *Resource Instance Service* (RIS). RIS 203 in the figure represent resource 207 as indicated by the dashed arrow 205 connecting them. Relationships and dependencies between resources in the physical infrastructure are represented, in the knowledge subsystem, by a canonical set of relationships between the corresponding RISs. The relationship 211 between RISs 203 and 213 indicate an inter-dependency between the corresponding 207 and 209 resources. The nature of this interdependency can be inferred by data associated with the relationship (such as type of relationship), not shown in the figure. Composite resources can also be represented and controlled by a RIS. For example, RIS 217 represent a composite resource which is a service environment (SE). A relationship of type *federates* connects RIS 217 with the three RISs 203, 213, and 215, as the arrows indicate. This relationship means that the three corresponding resources in the physical infrastructure are allocated to the service environment that RIS 217 represents. *Resource Managers (RMs)* represent and manage collections of resources of the same type. In particular, they provide an operation that returns a handle to a free resource instance. RM 219 manages the set of processing capacity resources. RMs 221, 223, and 241 manage the

1 collections of storage firewall, and software licenses resources, respectively. A composite
2 resource type may also have an RM that encapsulates the knowledge of how to build the
3 composite resource from other resources. An RM for a composite resource may use the RMs of
4 the resources that comprise the composite resource. For example, in the figure, RM 225 is an
5 RM for a *secure storage* composite resource; it uses RMs 221 and 223, as the arrows indicate.
6 Together, RMs and RISs encapsulate knowledge of resources capabilities, how resources can be
7 changed, how they can be associated with other resources, and what resources are currently free
8 or allocated. In managing resources, RMs and RISs implement methods that affect changes on
9 the resources they manage. These methods may be implemented in any way such as by scripts or
10 by a provisioning engine.

11 The management subsystem 239 includes a *Concrete Model Processing Engine (CMPE)* 227
12 which receives requests in the form of a Concrete Resource Model that describes a desired
13 resource structure. The CMPE generates provisioning actions 233 to reach a state that satisfies
14 the requirements specified in the Concrete Resource Model. Once these provisioning actions are
15 executed, either by the CMPE, or by the CMPE using a provisioning engine, they affect the state
16 of the system. For example, a request may describe a new resource structure for an existing or
17 non-existing service environment. Processing of a request results in provisioning actions that
18 may create a service environment or change the combination of resources allocated to a service
19 environment. In effect, the provisioning actions create a resource structure that matches the
20 description in the Concrete Resource Model. To generate the provisioning actions the CMPE
21 queries the knowledge subsystem, as indicated by arrow 235, to understand the state of the
22 system and how it can be changed. The CMPE generates provisioning actions that include
23 invocations of operations on knowledge subsystem entities, thus execution of a sequence of
24 provisioning actions affects the state of the system (resources in the physical infrastructure), only
25 through interaction with the knowledge subsystem (arrow 231), and not directly. The generation
26 and execution of provisioning actions may be interleaved; to serve a request, a sequence of
27 provisioning actions may be generated and executed before the next sequence is generated; a
28 sequence of provisioning actions may be regenerated if its execution fails.

1 It is important to note that the same method can be applied inside a RM for a composite resource.
2 Thus, the method can be applied to automatically generate a RM for a composite resource type
3 based on its definition using a Concrete Resource Model. The generated RM provides a set of
4 methods to create, destroy, or modify a composite resource based on a Concrete Resource Model
5 that describes its desired structure. The RM can be then used as any other RM by a higher level
6 CMPE. Such a strategy in fact distributes the CMPE method across infrastructure entities. It
7 also allows to easily create and reuse provisioning components in different level of granularity.

8 **Modeling Concepts**

9 Both the representation of the infrastructure using a Knowledge Subsystem and the description of
10 the desired service environment using a Concrete Resource Model rely on concepts taken from
11 the object-relationship paradigm. The Concrete Resource Model is an object-relationship model.
12 It can be implemented using UML, XML, relational databases, or similar technologies. A node
13 in this invention represents a resource of a group of resources. Relationships between objects in
14 the model represent relationships between resources. A set of relationship types is used to
15 represent the type of relationship. In particular, in the description of the invention we use the
16 following relationship types, defined below.

17 **A uses B:** if a resource **A** uses a resource **B** to perform its function.

18 **A hosts B:** if the resource **A** is a runtime container for **B**. For example: *Server host OS,*
19 *OS hosts application, WAS (WebSphere Application Server) hosts Enterprise Java Bean.*

20 **A connects B:** if **A** and **B** are connected (i.e., communication between **A** and **B** is
21 possible).

22 **A federates B:** if **A** is a resource that is a federation of other resources including **B**.

23 Namely, **A** represents a set of heterogeneous resources that cooperate in order to perform
24 a higher level function. For example, if a Web-site is a federation of Web-servers,
25 databases, and load-balancers, the following relationships are implied: *Web-Site federates*
26 *Web-server, Web-site federates database,* etc. Note that a composite resource is actually
27 modeled as a resource that *federates* the set of resources comprising it.

1 **A contains B:** if A is essentially a set of homogeneous resources. For example, *VLAN*
2 *includes a Switch Port.*

3 The invention does not assume any particular relationship type set. In particular, any relationship
4 type set can be used in a model. For the invention to work, each relationship type will have to be
5 mapped to an automation procedure in the knowledge subsystem for establishing or
6 un-establishing it. This mapping will be discussed later.

7 **The Knowledge Subsystem**

8 An knowledge subsystem describes the resources and the organization of resources in the service
9 provider's infrastructure. It includes the resource types and capabilities and information on
10 physical connectivity, such as the number of network adapters and the position of firewalls.

11 Infrastructure knowledge is distributed across entities of the knowledge subsystem, such as
12 RMs and RISs. In order to obtain knowledge necessary for its function the CMPE may query
13 entities of the knowledge subsystem.

14 Figure 3 describes infrastructure knowledge and the way it might be distributed among
15 knowledge subsystem entities. In the knowledge subsystem (619) RISs 603 and 605 represent
16 resource instances 607 and 609, respectively. Relationships between RISs represent operational
17 dependencies between the corresponding resources. For example, the relationship 601 of type
18 *uses* between RIS 603 representing server 607 and RIS 605 representing shared file system (SFS)
19 609 indicates that the server 607 uses the shared file system 609. A RIS can be queried for its set
20 of relationships, as well as for values of configuration attributes of the resource it represents.

21 RMs manage collections of resources of the same type. RMs can be queried to obtain
22 information about the collection, and to obtain information on a type level. For example, server
23 RM 611 encapsulates a server type model 613 that can be queried for possible relationships of a
24 server resource to other resources. The set of possible server relationships will include the
25 aforementioned *uses* relationship between a server and a shared file system, as well as other
26 relationships that the server may have. The knowledge subsystem also includes information on

1 infrastructure constraints and capabilities. Infrastructure constraints can be expressed using
2 rules, assertions or other mechanisms. We also use attributes on relationships (on a type or
3 instance basis) to express constraints.

4 Following are some mechanisms that can be used to describe such infrastructure constraints. A
5 *fixed* attribute on a relationship expresses that the relationship between the two corresponding
6 resources cannot be changed by a provisioning action. For example, in a wire-and-forget
7 environment, where resources are wired exactly once to a set of switches, and the wiring cannot
8 be changed, a *connects* relationship between a Network Interface Card (NIC) and a Switch Port
9 (SP) will have the *fixed* attribute.

10 **The Concrete Resource Model**

11 The *Concrete Resource Model* is the input to the CMPE process. It declaratively describes a
12 desired resource structure which may include new resources added to an existing or new service
13 environment, or resources removed from it. The goal of the CMPE process is to generate such a
14 resource structure on the computing utility infrastructure.

15 In a Concrete Resource Model, nodes represent resources, and requirements on the state of these
16 resources. Edges represent relationships between resources. Every edge is associated with a list
17 of attributes that describes the nature of the relationship. A node may include a set of constraints
18 on values of attributes of the resource that it represents. The values of some of these attributes
19 are *fixed*, namely, they cannot be changed. Therefore, the constraints on these attributes are used
20 as selection criteria for a resource that will serve the role of this node in the final resource
21 structure that implements the Concrete Resource Model in the end of the provisioning process.

22 As in the knowledge subsystem, relationships between nodes in the Concrete Resource Model
23 may be *fixed* or *dynamic*. *Fixed* relationships cannot be changed; they reflect fixed infrastructure
24 structures and operational constraints. Thus, such relationships must be taken into account in the
25 selection of the resources. For example, if a Concrete Resource Model includes a server node

1 with a *fixed contains* relationship with three NICs then only a server with (at least) three NICs
2 can be selected for this node. A *dynamic* relationship can be established by invoking a low level
3 automation procedure on one (or more) of the knowledge subsystem entities. For example a
4 dynamic *connects* relationship may be established between switch port and VLAN resources by
5 programmatically configuring switches or routers.

6 In the preferred embodiment an edge, representing a relationship, is associated with a set of
7 attributes that describe the nature of the relationship. Attributes describe the type of relationship
8 (e.g., *federates*), and whether it is *fixed* or *dynamic*. A *color* attribute with value *green* denotes
9 that the relationship must exist between the corresponding resources. The same attribute with
10 value *red* denotes that the relationship must not exist between the corresponding resources.
11 Clearly, the nature of the relationship, and assertions on whether it should be present or absent in
12 the resulting resource structure, can be expressed in many different ways, the *color* attribute is
13 just one way to do that.

14 A Concrete Resource Model is *fully implementable* on a given infrastructure represented by a
15 knowledge subsystem if it is *mappable* onto the knowledge subsystem. More specifically, every
16 node in the Concrete Resource Model that represents a resource has to be mappable, directly or
17 indirectly, to either an RM or an RIS. Every relationship has to be mappable to an automation
18 procedure to establish it (or unestablish it). The mapping might be indirect; if one resource, say a
19 server, has a *fixed* relationship with a different resource, say a NIC, then only the server node in
20 the Concrete Resource Model needs to be mappable to a server RM. During the CMPE process,
21 a server RIS can be obtained from the server RM, and a NIC RIS can then be obtained from the
22 server RIS. Note that the NIC is represented in the server type model encapsulated by the server
23 RM. This condition will be further explained when discussing the operation of the CMPE. If
24 only parts of the Concrete Resource Model can be mapped to the knowledge subsystem then the
25 method of the invention can still be applied to create a resource structure than matches in parts
26 with the Concrete Resource Model.

27 Figure 4 shows an example of a Concrete Resource Model for a scientific computing cluster

1 service described in Figure 5. In this system every server (depicted in Figure 5 as a box) is
2 connected to a designated Admin VLAN 801. A free server 813 has all of its other NICs
3 connected to a designated Free-pool VLAN 803. A service environment 813, is a scientific
4 computing cluster environment which includes a single master server 809, connected to both a
5 front end VLAN 805 and a back end VLAN 807, and a set of worker servers connected to the
6 back end VLAN. Both Master and Worker servers are also connected to the Admin VLAN.

7 In Figure 4, the root node 701 is an object that represents the service environment itself; it
8 *federates* four resources: a master node 703, a set of zero or more worker nodes 707, a front end
9 VLAN 711, and a back end VLAN 709. The master node *includes* three network adapters
10 (NICs): one 715 connecting it through a switch port 727 to the Admin VLAN 729, one 705
11 connecting it through a switch port 723 to the front end VLAN 711, and one 713 connecting it
12 through a switch port 725 to the back end VLAN 709. The worker node *includes* two network
13 adapters: one 717 connecting it through a switch port 731 to the admin VLAN, and one 719
14 connecting it through a switch port 721 to the back end VLAN. VLANs group switch ports;
15 Each VLAN is represented by a node which *includes* one or more switch ports.

16 The relationships between NICs and the servers and the NICs and the switch ports are *fixed*
17 relationships. They are defined when the example infrastructure is set up and this physical
18 connection is considered to be permanent. Such is also the case for the *contains* relationship
19 between the Admin NIC and a set of switch ports, as the servers in this example are to remain
20 permanently connected to the Admin VLAN. The dynamic relationships need to be established
21 by the CMPE after the resources, represented by the nodes in the Concrete Resource Model, are
22 selected. An example of a dynamic relationship is the relationship between the Back End VLAN
23 709 and the switch ports that it *includes* (725 and 721). This relationship is established by
24 programmatically re-configuring the switches. The CMPE identifies such tasks and invokes the
25 corresponding procedures in the knowledge subsystem to carry them out.

26 **CMPE Operation**

Figure 6 is an example of an execution of the CMPE which illustrates the process of performing provisioning based on a Concrete Resource Model. The input to the process is a Concrete Resource Model 907. During the execution of the process, resources corresponding to nodes in the Concrete Resource Model are selected or created and relationships between them are established by interacting with the RMs in the knowledge subsystem. There may be multiple phases in which resources are selected or created and then configured to establish the corresponding relationships. In this example, after step 909 three resources are selected; the corresponding nodes are shown in black 917. In step 911 two relationships between the selected resources are established as depicted in 919. In step 913 the rest of the resources are selected as shown in 921. Finally, in step 915 all (dynamic) relationships between the resources are established as shown in 923. After the termination of the process, a structure matching the Concrete Resource Model structure is created in the knowledge subsystem. Hereafter we describe the process in more detail.

The CMPE generates and executes provisioning actions to create a resource structure that matches the Concrete Resource Model. A provisioning action sequence includes two types of provisioning actions: an action to select a resource, and an action to configure a resource or relationship. Usually a sequence corresponding to a single request will include multiple sub-sequences, termed *phases*, in which resources are selected and then configured. The number of phases depends on the complexity of the system.

A single phase of the CMPE process is described in Figure 7. Essentially, in a single phase, starting in step 1201, a set of infrastructure resources are selected and mapped to nodes in the Concrete Resource Model (1203). This matching defines a set of provisioning actions of two types: a provisioning action to configure a resource, and a provisioning action to establish or unestablish a dynamic relationship. A provisioning action may have preconditions. A provisioning action can be executed only if its preconditions are satisfied. In step 1209 a provisioning action whose preconditions are satisfied is executed. An execution of a provisioning action may have side effect in the physical infrastructure. In step 1211 these side effects are reflected back in the knowledge subsystem by creating or destroying the

1 corresponding relationships or changing the values of attributes in the corresponding Resource
2 Instance Services (side effects will be further explained later.)

3 If all provisioning actions are executed successfully (1213) then the current phase of the CMPE
4 process terminates successfully (1223). If all nodes of the Concrete Resource Model were
5 matched in this or previous phases then the process terminates successfully and the entire
6 structure described in the Concrete Resource Model is now built in the physical infrastructure
7 and reflected in the knowledge subsystem. Otherwise, another phase is performed.

8 In some cases, a provisioning action whose preconditions are not satisfied exists (1213). In this
9 case, a different action to satisfy an unsatisfied precondition, if exists, is added to the set of
10 provisioning actions that need to be executed (1221). If actions exist whose preconditions are not
11 satisfied, and there do not exist any precondition that can be satisfied by executing a different
12 action (1219 and 1225) then the process fails (1227). In general, it may be the case that an action
13 to satisfy a precondition also has preconditions. The method is applied recursively and may lead
14 to a chain of actions that are executed in order to satisfy the precondition of the original
15 provisioning action, this case is taken care of by adding an action to satisfy a precondition to the
16 set of provisioning actions (1221).

17 We now describe some aspects of the process in more details. In the matching step, selection of
18 resources is based on two conditions: when a node is matched with a corresponding Resource
19 Instance Service the values of fixed attributes as defined in the Resource Instance Service must
20 satisfy constraints on these attributes defined in the node. In addition all edges that represent
21 fixed relationships with the node as an endpoint in the Concrete Resource Model must match the
22 set of fixed relationships of the matching Resource Instance Service in terms of type of
23 relationship, direction of relationships, and matching endpoint nodes. Other criteria, such as
24 multiplicity, if defined, can also be matched against. For example, if a node **A** in the Concrete
25 Resource Model is connected by an edge to a node **B** in the Concrete Resource Model and
26 annotates with type *x* and color green then there has to be a relationship of type *x* between the
27 Resource Instance Service **A'** that is matched with the node **A** and the Resource Instance Service

1 **B'** that is matched with the node **B**, moreover the direction of the relationship should be identical
2 to the direction of the aforementioned edge. If the color of the aforementioned edge is red, then
3 there must not be such a relationship between **A'** and **B'** in the knowledge subsystem. The
4 matching algorithms works by interacting with the RMs. For every node, an operation is
5 executed on the corresponding RM to find and obtain a set of Resources Instance Services that
6 are potential match for the node. The RMs may accept some selection criteria (in the form of
7 constraints over values of attributes) which are defined in the node and passed as parameters to a
8 *find* operation. These selection criteria only serve for the initial filtering. Additional filtering
9 must be done by the CMPE so that the matching condition defined above is satisfied. The
10 matching is intricate since when selecting a resource (represented by a Resource Instance
11 Service) it is not enough to look only at its immediate fixed relationships; a selection of a
12 resource may dictate selection of a different resource (with whom it has a relationship) so the
13 latter one must also have the correct set of fixed relationships recursively. To do the matching,
14 the CMPE employs well known graph matching techniques. These techniques backtrack and try
15 the next possibility whenever a matching possibility fails.

16 Once the matching is complete, a subset of the nodes is mapped to Resource Instance Services,
17 such that the set of fixed relationships and attributes in the Concrete Resource Model matches the
18 corresponding relationships and attributes in the knowledge subsystem. This includes fixed
19 relationships between nodes that are both matched in this phase, or between nodes one of which
20 is matched in this phase and the other was matched in previous phases. In 1203, if a nontrivial
21 matching (i.e., matching of size greater than 0) exists, then the process proceeds to configuring
22 the resources, starting at 1207, otherwise the process fails in 1205. A process may fail due to
23 many reasons. For example, it may not be possible to map the pattern defined by the set of fixed
24 relationships to the infrastructure at hand. For example, if a node representing a server has 3
25 fixed *contains* relationships with nodes representing NICs and all servers in the infrastructure
26 have fewer than 3 NICs.

27 If a matching cannot be found the entire process fails (1205). Since some allocation and
28 configuration actions may already have been performed, a compensation action needs to take

1 place to restore the system state (1229). This may be done by recursively calling the CMPE with
2 a new request such that the new desired state is the original state before the current CMPE
3 process started, other alternatives, such as keeping a log and rolling back, may also be used.
4 Alternatively, the algorithm can be easily generalized to find a non-optimal solution in which a
5 resource structure similar but not identical to the description in the Concrete Resource Model is
6 found.

7 Once Resource Instance Services are selected and mapped to a subset of the nodes in the model,
8 they are configured to establish the set of dynamic relationships described in the model and to
9 change values of dynamic attributes to satisfy the constraints defined in the Concrete Resource
10 Model. This is done by interacting with knowledge subsystem entities that encapsulate the logic
11 to configure the resources. Different knowledge subsystem entities may encapsulate automation
12 procedures to establish (or unestablish) different relationships. For example, a RM for a
13 composite resource may encapsulate the knowledge to establish all relationships between
14 resources in the composite. The invention does not make any assumption on the architectural
15 location of these automation procedures. It only assumes that such low level automation
16 procedures exist, and that there exists a mapping, accessible to the CMPE, between a relationship
17 and the automation procedure to establish or unestablish it, and between an attribute and the
18 operation to set its value. An automation procedure will typically receive as parameters the
19 handles for resources involved and configure them to implement the semantics of the relationship
20 or attribute. The CMPE is responsible for updating the corresponding Resource Instance
21 Services with the information on the established relationship or value of attribute.

22 As explained above, a matching of a set of Concrete Resource Model nodes and a set of
23 Resource Instance Services defines a set of provisioning actions that must be executed. The set
24 includes two types of provisioning actions: to configure a node and to establish or unestablish a
25 relationship. Specifically, for every node in the Concrete Resource Model, for every attribute
26 whose value is different then the value of the attribute in the matching Resource Instance Service,
27 a provisioning action must be performed on the Resource Instance Service to change the value of
28 the attribute. For every edge representing a dynamic relationship in the Concrete Resource

1 Model, a provisioning action must be performed to establish the relationship in the knowledge
2 subsystem if the relationship do not exist. For every red edge in the Concrete Resource Model, if
3 a corresponding relationship exists in the knowledge subsystem it must be un-established.

4 For example, a resource, say of type Web server, may have an attribute *state* whose value in the
5 Concrete Resource Model is set to *started*. In the knowledge subsystem the values of the *state*
6 attribute may be *created*. A provisioning action must be performed to change the value from
7 *created* to *started*. The process can work as follows. An operation *setAttribute* is invoked on the
8 Resource Instance Service with parameters that are the name of attribute and the new value
9 (“state”, and, “started”, correspondingly). This operation triggers an automation procedure which
10 affects the physical infrastructure by starting the Web server that is represented by the Resource
11 Instance Service. For a relationships, consider as an example a relationship of type *contains*
12 between a Switch Port and a VLAN. An *establishRelationship* provisioning action can be
13 invoked in the knowledge subsystem that will trigger an automation procedure that affects the
14 physical infrastructure by programmatically configuring the switch to move the designated
15 switch port in the designated VLAN.

16 Back in Figure 7, the configuration is a 3 step process; first, in 1207 all dynamic relationships
17 and attributes between matched resources are collected by analyzing the Concrete Model and
18 mapped to the corresponding provisioning action, second, in 1209, a provisioning action whose
19 preconditions are satisfied is executed.

20 As mentioned, a provisioning action may have side effects that need to be reflected back in the
21 knowledge subsystem. A side effect is any state change that is beyond the property that is the
22 target for which the automation procedure was invoked. An example of a side effect of the
23 operation to start a *Webserver*, described above, may be the creation of a *use* relationship
24 between the *Webserver* and a *database* resource. In this case, step 1211 involves updating the
25 knowledge subsystem with the aforementioned use relationships. Side effects may be modeled
26 and described in an inspectable way for every operation in the knowledge subsystem that affects
27 the physical infrastructure. Alternatively, they may be discovered by a different *discovery*

1 *component* after the operation is executed. In the later case once a provisioning action is
2 executed, the discovery component is executed and its output is used to update the knowledge
3 subsystem in Step 1211. In some settings the users of the system may decide that some
4 relationships are not important for the management of the system and they can be ignored
5 altogether.

6 Although the present invention may be employed by many types of entities, it is particularly
7 useful for use by a service provider, an enterprise owning an infrastructure used for running at
8 least one application, a customer of a service provider, a company owning an IT infrastructure,
9 and a utility provider.

10 The present invention includes a method for provisioning in a computing utility infrastructure.
11 The method comprising: obtaining a Concrete Resource Model describing a desired resource
12 structure; and using the Concrete Resource Model to generate at least one provisioning action to
13 create a matching resource structure in the computing utility infrastructure.

14 In some embodiments of the method: the step of using the Concrete Resource Model includes
15 executing at least one phase comprised of a matching step and a configuring step, wherein the
16 steps of matching and configuring are repeated until the Concrete Resource Model is entirely
17 matched with a set of at least one Resource Instance Service in said knowledge subsystem;
18 and/or the step of matching includes mapping every node in a subset of nodes in the Concrete
19 Resource Model to a Resource Instance Service in the knowledge subsystem, such that for every
20 node in said subset, constraints on values of fixed attributes in said node are satisfied by the
21 values of the same attributes in the corresponding Resource Instance Service, and the set of fixed
22 relationships between matched nodes matches the set of relationships between the corresponding
23 Resource Instance Services; and/or the step of configuring includes: selecting a provisioning
24 action having all preconditions satisfied, executing the provisioning action, and updating the
25 knowledge subsystem with side effects of the provisioning action, and repeating the steps of
26 selecting and executing until all provisioning actions whose preconditions are satisfied are
27 executed; and/or the step of configuring includes selecting a provisioning action having at least

1 one precondition not satisfied, adding said different action to satisfy said at least one
2 precondition to a provisioning action set to be executed; and/or further comprises obtaining said
3 side effects of an action by at least one of: inspecting the definition of said provisioning action in
4 said knowledge subsystem, and dynamically discovering the side effects once the action is
5 executed by executing a discovery component.

6 In further embodiments the step of provisioning includes at least one task taken from a group of
7 tasks consisting of: creating a new service environment; changing a combination of resources
8 allocated to the service environment; changing the configuration of resources allocated to a
9 service environment; destroying a service environment ; and any combination of these tasks.

10 In some embodiments the step of changing the configuration of resources allocated to a service
11 environment includes at least one of: changing a local state of a resource; and changing a way the
12 resource is configured to work with other resources.

13 In some further embodiments, the method further comprises regenerating provisioning
14 instructions whenever at least one of the following occurs: infrastructure characteristics change;
15 and requirements of the service change; and/or employing said Concrete Resource Model to
16 generate a Resource Manager for a composite resource

17 In some embodiments of the method, infrastructure characteristics are taken from a group of
18 characteristics consisting of: types of resources in the infrastructure; capabilities of said
19 resources; configuration of said resources; constraints on a configuration of said resources; and
20 any combination of these characteristics.

21 In some embodiments the method is employed to automatically generate a Resource Manager for
22 a composite resource; the Resource Manager provides a set of resource manager methods taken
23 from a group of resource manager methods consisting of: creating composite resources based on
24 a Concrete Resource Model; changing composite resources based on a Concrete Resource
25 Model; destroying composite resources based on a Concrete Resource Model; and any

1 combination of these resource manager methods.

2 In some embodiments of the method, the execution specification is taken from a group of
3 execution specifications consisting of: execute immediately; stored and executed later at least
4 once; and a combination of the above.

5 In some embodiments of the invention is an apparatus for provisioning in a computing utility
6 infrastructure, the apparatus comprising: means for obtaining a Concrete Resource Model
7 describing a desired resource structure; and means for using the Concrete Resource Model to
8 generate at least one provisioning action to create a matching resource structure in the computing
9 utility infrastructure.

10 Variations described for the present invention can be realized in any combination desirable for
11 each particular application. Thus particular limitations, and/or embodiment enhancements
12 described herein, which may have particular advantages to a particular application need not be
13 used for all applications. Also, not all limitations need be implemented in methods, systems
14 and/or apparatus including one or more concepts of the present invention.

15 The present invention can be realized in hardware, software, or a combination of hardware and
16 software. A visualization tool according to the present invention can be realized in a centralized
17 fashion in one computer system, or in a distributed fashion where different elements are spread
18 across several interconnected computer systems. Any kind of computer system - or other
19 apparatus adapted for carrying out the methods and/or functions described herein - is suitable. A
20 typical combination of hardware and software could be a general purpose computer system with
21 a computer program that, when being loaded and executed, controls the computer system such
22 that it carries out the methods described herein. The present invention can also be embedded in a
23 computer program product, which comprises all the features enabling the implementation of the
24 methods described herein, and which - when loaded in a computer system - is able to carry out
25 these methods.

1 Computer program means or computer program in the present context include any expression, in
2 any language, code or notation, of a set of instructions intended to cause a system having an
3 information processing capability to perform a particular function either directly or after
4 conversion to another language, code or notation, and/or reproduction in a different material
5 form.

6 Thus the invention includes an article of manufacture which comprises a computer usable
7 medium having computer readable program code means embodied therein for causing a function
8 described above. The computer readable program code means in the article of manufacture
9 comprises computer readable program code means for causing a computer to effect the steps of a
10 method of this invention. Similarly, the present invention may be implemented as a computer
11 program product comprising a computer usable medium having computer readable program code
12 means embodied therein for causing a a function described above. The computer readable
13 program code means in the computer program product comprising computer readable program
14 code means for causing a computer to effect one or more functions of this invention.

15 Furthermore, the present invention may be implemented as a program storage device readable by
16 machine, tangibly embodying a program of instructions executable by the machine to perform
17 method steps for causing one or more functions of this invention.

18 It is noted that the foregoing has outlined some of the more pertinent objects and embodiments of
19 the present invention. This invention may be used for many applications. Thus, although the
20 description is made for particular arrangements and methods, the intent and concept of the
21 invention is suitable and applicable to other arrangements and applications. It will be clear to
22 those skilled in the art that modifications to the disclosed embodiments can be effected without
23 departing from the spirit and scope of the invention. The described embodiments ought to be
24 construed to be merely illustrative of some of the more prominent features and applications of the
25 invention. Other beneficial results can be realized by applying the disclosed invention in a
26 different manner or modifying the invention in ways known to those familiar with the art.